2011 50th IEEE Conference on Decision and Control and
European Control Conference (CDC-ECC)
Orlando, FL, USA, December 12-15, 2011

The Intertemporal Utility of Demand and Price Elasticity of Consumption in Power Grids with Shiftable Loads

Mardavij Roozbehani † Ali Faghih ‡ Mesrob I. Ohannessian ‡ Munther A. Dahleh ♯

***Abstract*—This paper presents a mathematical model of consumer behavior in response to stochastically-varying electricity prices, and a characterization of price elasticity of consumption induced by optimally shifting flexible demands within a fixed time window. The approach is based on deriving the optimal load-shifting policy through a finite horizon stochastic dynamic program, and the analysis is performed under both perfect and partial information about price distribution. An aggregate demand model is constructed from individual demands with random arrivals and random deadlines. Under this model, the aggregate demand becomes a function of price only, and thus allows for quantitative characterization of the utility of demand and price elasticity. While the demand for electricity is often deemed to be highly inelastic, it is shown in this paper that optimal load-shifting can create a considerable amount of price elasticity, even when the aggregate consumption over a long period remains constant.**

***Index Terms*—Intertemporal Utility, Mathematical Modeling of Consumer Behavior, Price Elasticity, Dynamic Programming.**

I. Introduction

The advent of new generation, storage, and demand response technologies for improving energy efficiency will dramatically change the characteristics of supply and demand for electricity at both micro and macro levels. Consequently, econometric models of consumer behavior are extremely important for understanding the qualitative and quantitative characteristics of demand response and efficient integration of the related technologies. Motivated by these considerations, in this paper we present a mathematical model of consumer behavior in response to stochastically-varying electricity prices, and derive the associated optimal policies for shifting flexible loads under both partial and full information about price distribution. The closely related measure of responsiveness and sensitivity of demand to small variations in price, i.e., the *price elasticity* of demand is also examined within the same quantitative framework.

It has been asserted that at any given time, a considerable portion of the generated power is supplied to flexible loads that are shiftable in time [6]. Examples abound. Material processing, electric vehicle (EV) charging, heating, ventilation, air conditioning, refrigeration, and agricultural pumping can be shifted by a few minutes to a few hours at little or no cost. Some loads such as refrigeration, air-conditioning and heating can also be viewed as thermal storage via pre-cooling or pre-heating. Electric vehicles can be viewed as both electrical storage units and shiftable loads.

Previous research efforts concerning the effects of load-shifting and storage on price elasticity of demand have been mostly empirical and qualitative, see, e.g., [4]. In contrast, our approach is based on stochastic dynamic programming for optimal shifting of time-flexible loads in the presence of stochastically-varying prices. Related quantitative frameworks generally appear in the literature that address the consumer energy management problem, and are mostly based on dynamic programming.

While some earlier works such as [3] and [1] have introduced and discussed the general concepts of optimization of energy consumption under real-time electricity pricing, recent studies more related to our work have been reported in [5], [2], and [6]. In contrast to these studies, our model is more abstract, which enables us to obtain analytical and semi-analytical expressions for the solution of the underlying dynamic programming problem. These technical results provide an abstract description of a complicated system, and enable us to develop tractable models that effectively highlight the essential structural features of consumer behavior, most importantly in aggregate, from the grid operator's point of view. The more closely related paper by Papavasiliou and Oren [6] proposes a direct coupling of renewable generation with shiftable loads to mitigate the imbalances caused by the fluctuation of renewable energy supply. Their approach too, is based on stochastic dynamic programming. In particular, our result on the affine structure of the value function associated with the optimal load-shifting problem were partially obtained in [6].

The contributions of this paper are summarized as follows:

- We propose a mathematical model of consumer behavior in response to time-varying electricity prices under both perfect and partial information about price distribution (Sections II and III).
- We examine the relation between the economic value of load-shifting and price volatility. We present upper and lower bounds on the economic value for general classes of distributions, and exact relationships for some simple distributions (Section IV).
- Finally, we present a model that represents the aggregate response of a large number of consumers, and use simulations to characterize the utility of the aggregate demand under certain distributions. It is shown that load-shifting can induce a considerable amount of price elasticity, even when the cumulative consumption over a long period remains insensitive to price variations (Section V).

† Mardavij Roozbehani is a research scientist at the Laboratory for Information and Decision Systems (LIDS), at the Massachusetts Institute of Technology (MIT), Cambridge, MA. E-mail: mardavij@mit.edu

‡Ali Faghih and Mesrob I. Ohannessian are graduate students in Electrical Engineering and Computer Science (EECS), working at LIDS, MIT. Emails: afaghih@mit.edu, mesrob@mit.edu

♯ Munther A. Dahleh is the Associate Head of the EECS Department at MIT. E-mail: dahleh@mit.edu

For brevity, we have excluded all the proofs. The proofs can be found in the full version of this paper which is available at [10].

II. A Dynamic Model of Consumer Behavior

In this section, we develop a dynamic model of consumer behavior in response to stochastically-varying price signals. We formulate the consumer's energy management problem as an inventory control problem over a finite horizon.

A. Basic Definitions

Demand is the amount of electricity that the consumer would need to consume by withdrawing from the main grid to fulfill all the standing needs. *Consumption* is the amount of electricity that the consumer withdraws from the grid. In the absence of load-shifting and storage, consumption would be equal to demand. *Net consumption* is the amount of electricity that the consumer withdraws from, or sells back to the main grid. In the absence of local storage and local generation, net consumption would be equal to consumption.

B. Components of the Model

1) Demand: We denote the consumer's total demand at time $k \in \{0, \cdots, N\}$ by d_k, and assume that it consists of two components:

$$d_k = d_k^f + d_k^s$$

where d_k^f is the firm component and d_k^s is the shiftable component. It is assumed that at time $k = 0$, both d_k^s and d_k^f are perfectly known[1] to the consumer for all periods $k \in \{0, ..., N\}$. The shiftable demand d_k^s can be satisfied at any time $t \in \{k, ..., N\}$, whereas the firm demand d_k^f must be satisfied at time k. Both d_k^s and d_k^f are assumed to be inelastic.

2) The Decisions: The decision set of the consumer is characterized by a triplet

$$(u_k, v_k^{\text{in}}, v_k^{\text{out}}) \in [0, \overline{u}] \times [0, \overline{v}^{\text{in}}] \times [0, \overline{v}^{\text{out}}] \tag{1}$$

where, u_k is the amount of electricity that, at time k, the consumer allocates to fulfilling some or all of the shiftable demands, and v_k^{in} and v_k^{out} are, respectively, the amount of electricity that the consumer stores in, or withdraws from the storage. The corresponding upper bounds ($\overline{v}^{\text{in}}$ and $\overline{v}^{\text{out}}$) represent the physical ramp constraints on storage, while $\overline{u}$ limits the amount of electricity that can be allocated to satisfying backlogged demand. Letting $v_k = v_k^{\text{in}} - v_k^{\text{out}} \in [-v_k^{\text{out}}, v_k^{\text{in}}]$, the net consumption is given by

$$\begin{aligned} y_k &= v_k^{\text{in}} - v_k^{\text{out}} + u_k + d_k^f & (2)\\ &= v_k + u_k + d_k^f. & (3) \end{aligned}$$

It is assumed that y_k is constrained as:

$$y_k \in [-\underline{y}, \overline{y}], \qquad \underline{y}, \overline{y} \in [0, \infty) \tag{4}$$

Thus, $\underline{y} = 0$ corresponds to the situation where selling electricity back to the grid is not allowed.

[1] Our results can be extended to the case with stochastic uncertainty in both components of the demand.

3) The Price: Denoted by λ_k, the price process is a non-negative exogenous Markovian process driven by an independently distributed random process $\mathbf{w}_k$, according to

$$\lambda_{k+1} = g_k\left(\lambda_k, w_k\right)$$

where the functions g_k and possibly the distributions of $\mathbf{w}_k$ are known for each k. It is assumed that at the beginning of each discrete time interval $[k, k+1]$, the random variable λ_k is materialized and revealed to the consumer. A specific scenario where this model is readily applicable is where the distributions of the prices for the next 24 hours are computed in the day-ahead market and made available to the consumer. In this case, we may choose $g\left(\lambda_k, w_k\right) = w_k$, where the distribution of w_k is known.

We also assume that the feed-in and usage tariffs are the same, i.e., λ_k is the price per unit for both consumption ($y_k \geq 0$) and production ($y_k \leq 0$), and there are no transaction costs.

4) The States: It is assumed that the state of the consumer is characterized by a pair

$$(x_k, s_k) \in (-\infty, 0] \times [0, \overline{s}] \tag{5}$$

where x_k represents the amount of backlogged demand, and s_k represents the amount of stored energy. The parameter $\overline{s} \geq 0$ is the physical upper bound on the capacity of storage. We impose a deadline on backlog by requiring $x_N = 0$.

The states x_k and s_k evolve according to:

$$x_{k+1} = x_k + u_k - d_k^s \tag{6}$$

$$s_{k+1} = \beta s_k + \eta^{\text{in}} v_k^{\text{in}} - \eta^{\text{out}} v_k^{\text{out}} \tag{7}$$

where, u_k, v_k^{in}, and v_k^{out} are defined as in (1), $\beta \leq 1$ is the decay factor, $\eta^{\text{in}} \leq 1$ and $\eta^{\text{out}} \geq 1$ are charging and discharging efficiency factors For an idealized model of the dynamics of storage we have: $\beta = 1$, $\eta^{\text{in}} = 1$, and $\eta^{\text{out}} = 1$.

5) Disutility and Penalty: In general, there is a disutility associated with backlogging the demand. We characterize this disutility via a sequence of *cost* functions $p_k(\cdot)$ which essentially represent, in an abstract sense, the inconvenience associated with deferring the currently standing demands. Similarly, the cost associated with storage is characterized via a sequence of penalty functions $h_k(\cdot)$.

6) The Optimization-Based Model: The consumer's energy management problem can be formulated as a finite-horizon dynamic programming problem as follows

$$\begin{aligned} \min \quad & \mathbf{E}\left[\sum_{k=0}^{N} p_k\left(x_k\right) + h_k\left(s_k\right) + \lambda_k y_k\right] & (8)\\ \text{s.t.} \quad & x_{k+1} = x_k + u_k - d_k^s, \qquad x_N = 0 \\ & s_{k+1} = \beta s_k + \eta^{\text{in}} v_k^{\text{in}} - \eta^{\text{out}} v_k^{\text{out}} \\ & \lambda_{k+1} = g\left(\lambda_k, w_k\right) \\ & y_k = u_k + v_k^{\text{in}} - v_k^{\text{out}} + d_k^f \end{aligned}$$

The optimization problem (8) is further subject to state and action space constraints (1), (4), (5).

We now establish conditions under which (8) can be decomposed into two sub-problems: 1. optimal load-shifting, and 2. optimal storage management, as defined below:

1) *The load-shifting problem*:

$$\begin{aligned}\min \quad & \mathbf{E}\left[\sum\nolimits_{k=0}^{N} p_k(x_k) + \lambda_k u_k\right] \qquad (9)\\ \text{s.t.} \quad & x_{k+1} = x_k + u_k - d_k^s, \qquad x_N = 0\\ & \lambda_{k+1} = g_k(\lambda_k, w_k)\\ & x_k \in (-\infty, 0]\\ & u_k \in [0, \overline{u})\end{aligned}$$

2) *The storage problem*:

$$\begin{aligned}\min \quad & \mathbf{E}\left[\sum\nolimits_{k=0}^{N} h_k(s_k) + \lambda_k v_k\right] \qquad (10)\\ \text{s.t.} \quad & s_{k+1} = s_k + v_k\\ & \lambda_{k+1} = g_k(\lambda_k, w_k)\\ & s_k \in [0, \overline{s}]\\ & v_k \in [-\overline{v}^{\text{out}}, \overline{v}^{\text{in}}]\end{aligned}$$

Proposition. *Let γ_c^*, γ_x^*, and γ_s^* be, respectively, the optimal solutions to the general optimization problem (8), the load-shifting problem (9), and the storage problem (10). Let γ_f be the expected cost of the firm demand:*

$$\gamma_f = \mathbf{E}\left[\sum\nolimits_{k=0}^{N} \lambda_k d_k^f\right]$$

Then, for sufficiently large $\overline{u}$ we have

$$\gamma_c^* \geq \gamma_x^* + \gamma_s^* + \gamma_f.$$

Furthermore, suppose that both $\overline{y}$ and $\overline{u}$ are sufficiently large, and that the storage is lossless, i.e., $\beta = 1, \eta^{in} = 1$, and $\eta^{out} = 1$. Then

$$\gamma_c^* = \gamma_x^* + \gamma_s^* + \gamma_f.$$

The implication of this proposition is that when the limits ($\overline{y}$ and $\overline{u}$) are sufficiently large, the storage is ideal, and the feed-in and usage tariffs are the same, then, the consumer is indifferent to satisfying the demand by withdrawing from the grid, or from the storage. Hence, not only the problems of storage and load-shifting can be solved separately, but also the firm demand d_k^f becomes irrelevant to decision-making. The rest of the paper is devoted to the load-shifting problem (9). The study of the storage problem is addressed in a separate publication [9].

III. Optimal Policies for Load-Shifting

A. Perfect Information about the Price Distribution

In this section, we derive the consumer's optimal policy for scheduling the shiftable loads using stochastic dynamic programming under full information about price distribution.

Definition. *Given a probability mass function (PMF) P with support over a discrete set $\Theta \subseteq [\lambda_{\min}, \lambda_{\max}] \subset \mathbb{R}_+$, and cumulative distribution function F, the modulated double cumulative distribution function (MDCDF) associated with P is a concave function $\Gamma_P : [\lambda_{\min}, \infty) \to \mathbb{R}_-$ defined according to*

$$\Gamma_P(x) = \sum_{\theta=\lambda_{\min}}^{\min(\lambda_{\max}, x)} (\theta - x) P(\theta) \qquad (11)$$

When the support of P over the set Θ is uniformly spaced with distance L, then it can be shown using summation by parts that

$$\Gamma_P(x) = -L \sum_{\theta < x} F(\theta)$$

Theorem 1. Characterization of the Optimal Policy. *Consider the load-shifting problem (9) with sufficiently large $\overline{u}$,*

$$\overline{u} \geq \sum\nolimits_{k=0}^{N} d_k^s,$$

and i.i.d. price process λ with PMF P, and linear disutilities of delay, i.e., $p_k(x_k) = -p_k x_k$. Then,

(i) The optimal policy is a threshold policy characterized by

$$u_k^* = \begin{cases} 0 & \text{if} \quad \lambda_k > t_k \\ d_k^s - x_k & \text{if} \quad \lambda_k \leq t_k \end{cases} \qquad (12)$$

where, the thresholds t_k can be computed via the following recursive equations:

$$t_N = \lambda_{\max}, \quad t_k = p_k + t_{k+1} + \Gamma_P(t_{k+1}) \qquad (13)$$

where $\Gamma_P(\cdot)$ is the MDCDF (11).

(ii) The value function is affine, of the form:

$$V_k(x_k) = -t_k x_k + e_k. \qquad (14)$$

where the thresholds t_k are given in (13) *and the constant terms e_k satisfy the following recursive equations:*

$$e_N = 0, \quad e_k = e_{k+1} + d_k^s(t_{k+1} + \Gamma_P(t_{k+1})) \qquad (15)$$

(iii) Given $x_0 = 0$, the optimal expected cost is a demand-weighted sum of the differences between the stage thresholds and the marginal disutilities, i.e.,

$$V_0(0) = \sum_{k=0}^{N-1} d_k^s (t_k - p_k) \qquad (16)$$

B. Partial Information about the Price Distribution

In this section, we relax the *perfect information* assumption and propose a robust approximation scheme for the case where only mean and variance of the price distribution is known. This, we accomplish by approximating the function Γ_P, which embodies the dependence of the optimal policy on price distribution.

Theorem 2. Bounding the MDCDF. *Let $\lambda_{\min} = 0$ and $\lambda_{\max} = 1$. Given a mean $\mu \in [0, 1]$ and an achievable variance σ^2, let $\mathcal{P}$ be the set of all distributions with support on $[0, 1]$ that have mean μ and variance σ^2, and let $P \in \mathcal{P}$. Then, Γ_P can be bounded from above and below as follows:*

$$\underline{\Gamma}(x) \leq \Gamma_P(x) \leq \overline{\Gamma}(x), \quad \forall x \in [0, 1],$$

where,

$$\overline{\Gamma}(x) = \begin{cases} 0 & ; \; x \in \left[0, \mu - \frac{\sigma^2}{1-\mu}\right] \\ (1-\mu)(\mu - x) - \sigma^2 & ; \; x \in \left[\mu - \frac{\sigma^2}{1-\mu}, \mu + \frac{\sigma^2}{\mu}\right] \\ \mu - x & ; \; x \in \left[\mu + \frac{\sigma^2}{\mu}, 1\right] \end{cases}$$

$$\underline{\Gamma}(x) = \begin{cases} \frac{-\sigma^2}{\sigma^2+\mu^2}x & ;\ x \in \left[0, \frac{\mu^2+\sigma^2}{2\mu}\right] \\ \frac{-\sigma^2\sqrt{(\mu-x)^2+\sigma^2}}{\sigma^2+(\mu-x+\sqrt{(\mu-x)^2+\sigma^2})^2} & ;\ x \in \left[\frac{\mu^2+\sigma^2}{2\mu}, \frac{1-\mu^2-\sigma^2}{2(1-\mu)}\right] \\ \frac{-(1-\mu)^2(x-1)}{(1-\mu)^2+\sigma^2} + \mu - 1 & ;\ x \in \left[\frac{1-\mu^2-\sigma^2}{2(1-\mu)}, 1\right] \end{cases}$$

Furthermore, both of these bounds are tight pointwise, in the sense that for every $x \in [0,1]$ *there exists a distribution* $\overline{P} \in \mathcal{P}$ *for which* $\Gamma_{\overline{P}}(x) = \overline{\Gamma}(x)$ *and another distribution* $\underline{P} \in \mathcal{P}$ *for which* $\Gamma_{\underline{P}}(x) = \underline{\Gamma}(x)$. *Finally, the bounds for arbitrary (nonnegative)* $\lambda_{\min}$ *and* $\lambda_{\max}$ *can be obtained via the following transformations:*

$$\overline{\Gamma}(x; \lambda_{\min}, \lambda_{\max}, \mu, \sigma) = l\overline{\Gamma}\left(x_l; 0, 1, \mu_l, \sigma_l\right), \tag{17}$$

$$\underline{\Gamma}(x; \lambda_{\min}, \lambda_{\max}, \mu, \sigma) = l\underline{\Gamma}\left(x_l; 0, 1, \mu_l, \sigma_l\right), \tag{18}$$

where $l = \lambda_{\max} - \lambda_{\min}$, $x_l = (x - \lambda_{\min})/l$, $\mu_l = (\mu - \lambda_{\min})/l$, *and* $\sigma_l = \sigma/l$.

The proof of this theorem, which we omit for brevity, relies on the fact that the pointwise bounds can be posed as linear programs, the solutions to which suggest optimal distributions that are mixtures of impulses.

Corollary 1. Min-Max Approximation of the MDCDF. *For* $x \in [\lambda_{\min}, \lambda_{\max}]$, *let*

$$\widehat{\Gamma}(x) = \arg\min_{\ell} \max_{P \in \mathcal{P}} |\Gamma_P(x) - \ell|.$$

Then,

$$\widehat{\Gamma}(x) = \frac{\underline{\Gamma}(x) + \overline{\Gamma}(x)}{2}. \tag{19}$$

To illustrate these approximations, consider the case where $\mathcal{P}$ represents distributions over $[0,1]$, with mean $1/2$ and variance $1/12$. In Figure 1, we plot Γ_P for the special case of a uniform distribution with these parameters, give the upper and lower bounds of Theorem 2 over $\mathcal{P}$, as well as the min-max approximation of equation (19).

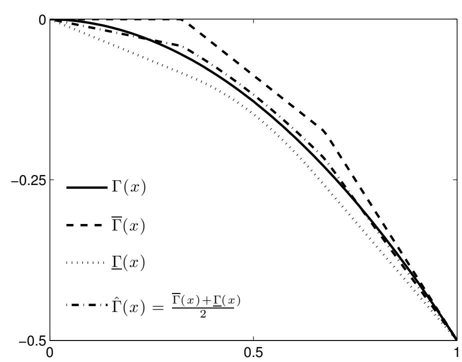

Fig. 1. For distributions over $[0,1]$, with mean $1/2$ and variance $1/12$, we illustrate the Γ_P of the uniform distribution, as well as partial information upper and lower bounds, and the min-max approximation.

The observation that the upper and lower bounds are not too far apart, suggests that partial information should not drastically change the characteristics of the consumer's response. We revisit this theme in the next two sections, where we study consumer savings and aggregate price elasticity.

We will use Theorem 3 and Corollary 1 to provide robust policies under partial information, and establish bounds on the optimal value under such policies.

Given a sequence $\{t_k\}$, we say that π is a threshold policy associated with $\{t_k\}$, if the control policy under π is of the form (12). We denote by $J_{P,\pi}$ the expected cost under threshold policy π, when the distribution function of λ is P.

Theorem 3. Bounding the Costs of Approximate Policies. *Let* $\mathcal{P}$ *be the set of all distributions with support on* $[\lambda_{\min}, \lambda_{\max}]$, *that have mean* μ *and variance* σ^2. *Let* $\overline{\Gamma}$, $\underline{\Gamma}$, *and* $\widehat{\Gamma}$ *be defined according to* (17)–(19). *Let* $\{\overline{t}_k\}$, $\{\underline{t}_k\}$, *and* $\{\hat{t}_k\}$, *be sequences recursively generated according to* (13) *with* $\overline{\Gamma}$, $\underline{\Gamma}$, *and* $\widehat{\Gamma}$ *respectively, and let* $\overline{\pi}$, $\underline{\pi}$, *and* $\hat{\pi}$, *be the associated threshold policies. Define:*

$$\overline{J} = \sum_{k=0}^{N-1} d_k^s \overline{t}_k, \qquad \underline{J} = \sum_{k=0}^{N-1} d_k^s \underline{t}_k, \qquad \widehat{J} = \sum_{k=0}^{N-1} d_k^s \hat{t}_k.$$

The following statements hold:

(i) *The threshold policy* $\overline{\pi}$ *is robust in the sense that*

$$\underline{J} \leq J_{P,\overline{\pi}} \leq \overline{J}, \quad \forall P \in \mathcal{P}.$$

(ii) *For any distribution* $P \in \mathcal{P}$, *the optimal cost is lower bounded by* $\underline{J}$. *That is,*

$$\underline{J} \leq J_P^* \stackrel{def}{=} \sup_{\pi} J_{P,\pi}, \quad \forall P \in \mathcal{P}.$$

(iii) *The expected cost under the min-max threshold policy* $\hat{\pi}$ *is close to* $\widehat{J}$ *in the following sense:*

$$\left|J_{P,\hat{\pi}} - \widehat{J}\right| \leq \sum_{k=1}^{N-1} f(k) \left|\Gamma_P\left(\hat{t}_k\right) - \widehat{\Gamma}\left(\hat{t}_k\right)\right|, \quad \forall P \in \mathcal{P}$$

$$\leq \frac{1}{2} \sum_{k=1}^{N-1} f(k) \left|\overline{\Gamma}\left(\hat{t}_k\right) - \underline{\Gamma}\left(\hat{t}_k\right)\right|,$$

where

$$f(k) = \sum_{i=0}^{k-1} d_k^s.$$

(iv) *The upper bound* $\overline{J}$ *is a quadratic function of* σ.

IV. The Value of Load-Shifting

We define the *value of load-shifting* as the expected savings that result from adoption of the optimal policy (12), when compared with consuming electricity on demand. The formal definition is as follows:

$$\mathcal{V} = \sum_k d_k^s \mu - \mathbf{E}[V_0(x_0)], \tag{20}$$

where $V_k(\cdot)$ is the value function defined in (14). Thus, we are benchmarking the performance of the optimal policy (12) against the average cost in the absence of load-shifting. Intuitively, we expect that the value of load-shifting would increase with price volatility, as measured by the variance. The following corollary establishes that as volatility of the price varies while the mean price is being held constant, the value of load-shifting is lower bounded by a quadratic function of the standard deviation for all possible distributions. The presented bounds are applicable to both partial and full information cases.

Corollary 2. *Let $\mathcal{V}$ be the value, as defined in* (20), *of the load-shifting problem (9). Then, for all distributions with fixed mean μ and support over a bounded subset $[\lambda_{\min}, \lambda_{\max}] \subset \mathbb{R}_+$, the following statements hold:*

(i) *There exists a constant C, such that for all achievable variances σ^2 we have*

$$\sum_{k=0}^{N-1} d_k^s(\mu - \underline{t}_k) \geq \mathcal{V} \geq C\sigma^2 = \sum_{k=0}^{N-1} d_k^s(\mu - \bar{t}_k)$$

where $\{\underline{t}_k\}$ and $\{\bar{t}_k\}$ are defined as in Theorem 3. Furthermore, regardless of the the information structure about the distribution, both bounds are tight for $N = 2$.

(ii) *Under full information, for all uniform distributions, $\mathcal{V}$ is a linear function of σ.*

We compute these bounds for the case where $\lambda_{\min} = 0$, $\lambda_{\max} = 100$, and the mean is $\mu = 50$. Figure 2 shows the upper and lower bounds on $\mathcal{V}$ as the standard deviation σ varies from 0 to the maximum achievable $\overline{\sigma} = 50$. Note that the lower bound applies to both worst-case distribution under full information, and to any distribution under policy $\overline{\pi}$ (cf. Theorem 3).

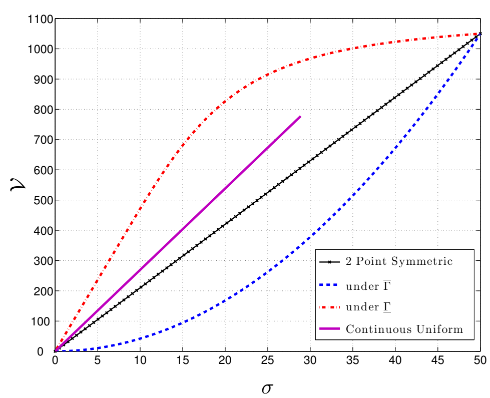

Fig. 2. Illustration of upper and lower bounds on the value of load-shifting as a function of the standard deviation, under all possible distributions with support over $[0, 100]$ and mean $\mu = 50$. The shifting horizon is $N = 24$. Also shown are the value under 2-point symmetric distribution and continuous uniform distribution with the same mean $\mu = 50$. Note that the maximum achievable standard deviation under continuous uniform distribution is $100/\sqrt{12} \simeq 28.86$.

V. Utility and Price Elasticity

In Section III we observed that the optimal response of an individual consumer to a given price is determined by both the thresholds, which are time-varying functions of the distribution, and the amount of backlogged demand, which in turn depends on the history of the stochastic price process. The response is therefore, a time-varying function of the internal state of the consumer. In this section we introduce an aggregation model that eliminates the dependencies of consumption on time and the internal states of the consumers, and thereby, allows us to characterize the collective behavior of a large number of consumers as a function of price only. The details of the model are as follows.

A. Aggregation Model

The aggregation model is specified as follows. We denote the total number of consumers by L, and assume that there is a global time horizon $\overline{T}$ for all consumers. At any given time $k \in \{0, \cdots, \overline{T}\}$, all consumers are given the same price signal λ_k. Each consumer $j \in \{1, \cdots, L\}$ starts consumption at a random time $T_{\text{start}}(j)$ that is uniformly distributed over $\{0, \cdots, \overline{T} - 1\}$, and has a random deadline $N(j)$ that conditioned on $T_{\text{start}}(j)$, is uniformly distributed over $\{1, \cdots, \overline{T} - T_{\text{start}}(j)\}$. Note that we could equivalently interpret these as different jobs, perhaps belonging to the same consumer, that arrive with a certain deadline in the system. The total consumption of all consumers at time k is the ensemble average of the individual u_k^j values, which for $k \in \{T_{\text{start}}(j), \cdots, T_{\text{start}}(j) + N(j)\}$ is determined by (12), and is taken to be zero if $k \notin \{T_{\text{start}}(j), \cdots, T_{\text{start}}(j) + N(j)\}$.

The dependence on state is eliminated by averaging over the process history. In particular, we can define:

$$\overline{u}^j(t, \lambda) = \mathbf{E}\left[u_t^{j*} | \lambda_t = \lambda\right].$$

In order to eliminate stage-dependence, we think of the consumption-measuring observer as sampling a random time τ uniformly over $\{0, \cdots, \overline{T}\}$. By averaging over this randomness, we maintain dependence on price alone. More precisely, we are interested in the quantity:

$$u_{\text{a}}(\lambda) = \frac{1}{L} \sum_{j=1}^{L} \mathbf{E}_\tau \left[\overline{u}^j(\tau, \lambda)\right],$$

which is easily captured in numerical simulations by clustering real-time prices, and averaging over each cluster.

B. Simulations

1) Aggregation Parameters: In the numerical simulations we used $L = 500$ (number of consumers), $\overline{T} = 720$ (global horizon). The choice of $\overline{T} = 720$ corresponds to a period of 24 hours with price updates in every two minutes.

2) Load-Shifting Model Parameters: We assume that each consumer starts with no backlogged demand, i.e. $x_0 = 0$, and that $d_k^s = 1$ for all k. In order to investigate how load-shifting penalties affect the aggregate utility, we examine two scenarios: *(a)* $p_k = 0$ for all k, and *(b)* $p_k = 0.1\mu$ for all k.

3) The Distribution Parameters: We simulate with three different price distributions: a discrete uniform distribution, a 3-point symmetric distribution with support in {0,0.5,1} and P(0.5)=0.5, and a discretized and truncated log-normal distribution, using the same mean $\mu = 0.5$ across distributions.

4) Numerical Results: Figure 3 illustrates how the normalized aggregate demand changes as a function of price for each of the three distributions in scenario (a), i.e., when $p_k = 0$ for all k. Each plot contains two graphs; one graph represents the aggregate consumption for the load-shifting problem (9) where consumers have perfect information about the price distribution and the other represents the same quantity where users substitute the true threshold function with the min-max optimal approximation (19) under partial information.

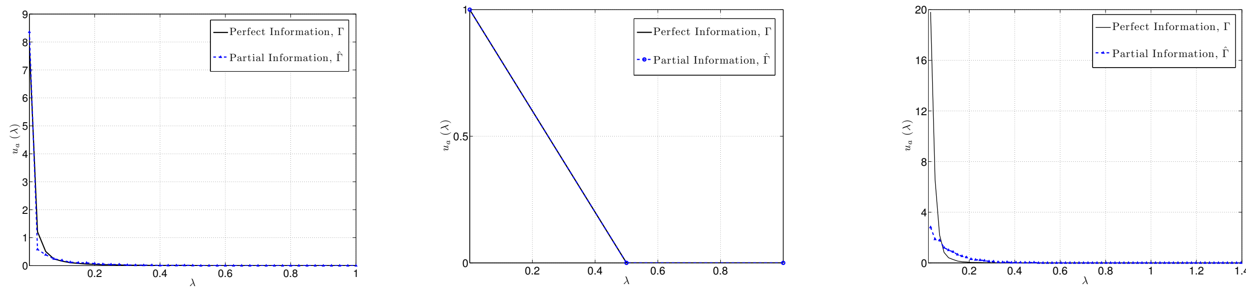

Fig. 3. Aggregate demand vs. price, with no backlog penalty, for the uniform (left), 3-point symmetric (middle), and truncated log-normal (right) distributions.

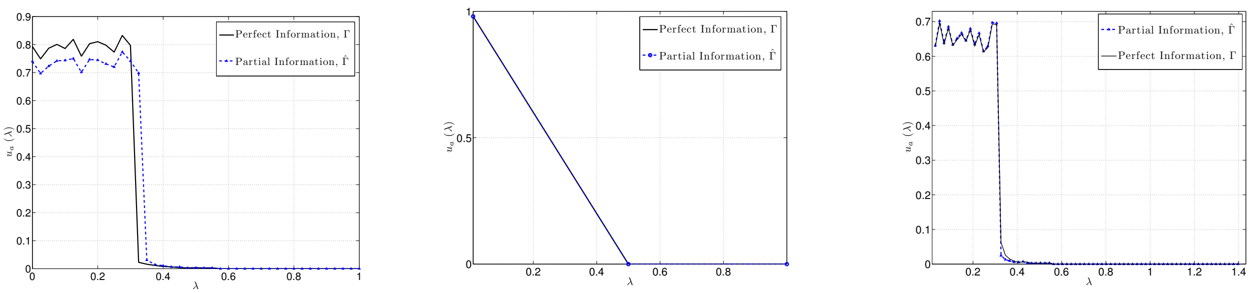

Fig. 4. Aggregate demand vs. price, with backlog penalty, for the uniform (left), 3-point symmetric (middle), and truncated log-normal (right) distributions.

Figure 4 illustrates how the aggregate demand changes as a function of price for each of the three distributions in scenario (b), i.e. when $p_k = 0.1\mu$ for all k. As before, each plot contains two graphs, corresponding to perfect or partial information about the price distribution.

C. Interpretation

In the absence of penalties (i.e. when $p_k = 0$ for all k, Fig. 3), the aggregate consumption varies as a relatively smooth function of price. However, when the demand has even a small disutility of delay, (10 precent of the average price in these simulations, i.e., $p_k = 0.1\mu$ for all k), the aggregate demand behaves approximately according to the policy of purchasing from the grid only when prices are below μ, see Fig. 4. This means that the price elasticity of demand is very small nearly everywhere, except when the price is close to a certain threshold near μ, where the demand shows significant elasticity. For the symmetric 3-point distribution, it follows from the above analysis that if at any given time x percent of the overall load is shiftable, ternary pricing could induce a price elasticity of nearly $100x/(1+x)$ percent at the middle and lowest prices, and zero at the highest price.

In view of the results of [7], [8], the implications of this observation is significant for stability of power grids under dynamic pricing.

VI. Conclusions and Future Work

We proposed a dynamic model of consumer behavior in response to stochastically-varying electricity prices and derived optimal policies for the load-shifting problem. We also proposed an approximation to the optimal policy for the case when the price distribution is not known, but only its support, mean, and variance are given. We showed that for all distributions with bounded support, the value of load-shifting, defined as the total expected savings from implementing the optimal load-shifting policies is at least of the same order as the variance of the price.

An interesting observation is that load-shifting can induce a considerable amount of price elasticity. The actual amount would depend on the ratio of shiftable loads to firm loads, which is time-varying and depends on the temporal characteristics of the load.

Future work includes extending these studies to analysis of the feedback interconnection of consumes and the wholesale markets and investigate the implications of load-shifting and storage on volatility and robustness of the closed loop system.

References

[1] J. W. Black, *Integrating Demand into the U.S. Electric Power System: Technical, Economic, and Regulatory Frameworks for Responsive Load*, Ph.D. Thesis, MIT, 2005.

[2] E. Bitar, A. Giani, R. Rajagopal, D. Varagnolo, P. Khargonekar, K. Poolla, P. Varaiya. "Optimal contracts for wind power producers in electricity markets." *Proc. of 49th IEEE Conf. on Decision and Control*, Atlanta, GA, 2010.

[3] P. Constantopoulos, F. C. Schweppe, and R. C. Larson. "A Real-Time Consumer Control Scheme for Space Conditioning Usage under Spot Electricity Pricing." *J. of Computers and Operations Research*, vol. 18, no. 8, pp. 751–765, 1991.

[4] D. S. Kirschen. "Demand-side view of electricity markets." *IEEE Trans. Power Syst.*, vol. 18, no. 2, pp. 520–527, May 2003.

[5] D. Livengood and R. Larson, "Energy Box: Locally Automated Optimal Control of Residential Electricity Usage." *Service Science*, vol. 1, pp. 1–16, 2009.

[6] A. Papavasiliou and S. S. Oren. "Supplying renewable energy to deferrable loads: algorithms and economic analysis." *IEEE PES General Meeting*, Minneapolis, Minnesota, July 25-29 2010.

[7] M. Roozbehani, M. A. Dahleh, and S. K. Mitter, "Volatility of power grids under real-time pricing," *Power Systems, IEEE Transactions on*, pp. Submitted, available at http://web.mit.edu/mardavij/www/publications.html, Mar. 2011.

[8] ——, "On the stability of wholesale electricity markets under real-time pricing," in *Decision and Control (CDC), 2010 49th IEEE Conference on*, Dec. 2010, pp. 1911 –1918.

[9] A. Faghih, M. Roozbehani, and M. A. Dahleh, "Optimal utilization of storage and the induced price elasticity of demand in the presence of ramp constraints" in *Decision and Control (CDC), 2011 50th IEEE Conference on*, Dec. 2011, To Appear.

[10] http://arxiv.org/ or http://web.mit.edu/mardavij/www/publications.html